STEGANALYSIS IN RESIZED IMAGES

Jan Kodovský *Jessica Fridrich*

Binghamton University
ECE Department
Binghamton, NY

ABSTRACT

It is well known that the security of a given steganographic algorithm strongly depends on the statistical properties of the cover source. In this paper, we study how downsampling affects steganographic security. The secure payload no longer scales according to the square-root law because resizing changes the statistical properties of the cover source. We demonstrate this experimentally for various types of resizing algorithms and their settings and thoroughly interpret the results. Modeling digital images as Markov chains allows us to compute the Fisher information rate for the simplest resizing algorithm with the box kernel and derive the proper scaling of the secure payload with resizing. The theory fits experimental data, which indicates the existence of a new scaling law expressing the length of secure payload when the cover length is not modified by adding or removing pixels but, instead, by subsampling. Since both steganography and steganalysis is today commonly evaluated through controlled experiments on resized images (e.g., the BOSSbase), the effect of resizing on security is of utmost importance to practitioners.

Index Terms— Steganalysis, resized images, interpolation algorithm, steganographic security

1. INTRODUCTION

Statistical properties of cover source strongly affect steganographic security. For instance, the Square-Root Law (SRL) [1] states that a constant level of statistical detectability is obtained when the message length grows proportionally to the square root of the number of pixels in the image. This law manifests when pixels *from the same source* are added/removed, e.g., by stitching together images to obtain a panorama or by cropping. When the number of pixels is changed by *resizing*, the statistical properties of the source change and the SRL no longer holds in its standard form. The investigation of this phenomenon is the subject of this paper.

This research direction is highly relevant to practitioners since the established method for benchmarking steganography as well as steganalysis is routinely done on resized images. In fact, today's most commonly used data set is the BOSSbase database [2] whose images were resized from their native resolution.

In the next section, we start with a simple motivational experiment that shows how much the statistical detectability of the HUGO algorithm [3] depends on the resizing kernel used to downsample the original full-resolution images forming the BOSSbase. To better isolate the effects of resampling, in Section 3 we work with a homogeneous source and show experimentally that resizing algorithms and their settings substantially change the security when measured empirically using detectors implemented as binary classifiers. In Section 4, we provide a theoretical analysis of the observed results for the nearest neighbor resizing by adopting a Markov chain model for the cover source. For this type of cover source and resizing algorithm, there exists a closed-form expression for the steganographic Fisher information rate, which allows us to determine the size of the secure payload that leads to the same level of statistical detectability. The paper is concluded in Section 5 where we discuss future directions and extensions of this work.

2. MOTIVATIONAL EXPERIMENT

The BOSSbase image data set has been introduced to the steganographic community during the BOSS competition [2]. The newest BOSSbase version 1.01 consists of 10,000 images obtained from full-resolution RAW images (coming from eight different cameras) by first demosaicking them using UFRaw (`http://ufraw.sourceforge.net/`), converting to 8-bit grayscale, resizing so that the smaller side is 512 pixels, and finally central-cropping to 512×512 pixels. All operations were carried out using the ImageMagick's `convert` command-line tool. The parameters of the UFRaw's demosaicking algorithm are not important for the purpose of this paper – they were fixed for all experiments to the same values as in the script used in the BOSS competition.

The work on this paper was supported by Air Force Office of Scientific Research under the research grant number FA9950-12-1-0124. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright notation there on. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied of AFOSR or the U.S. Government.

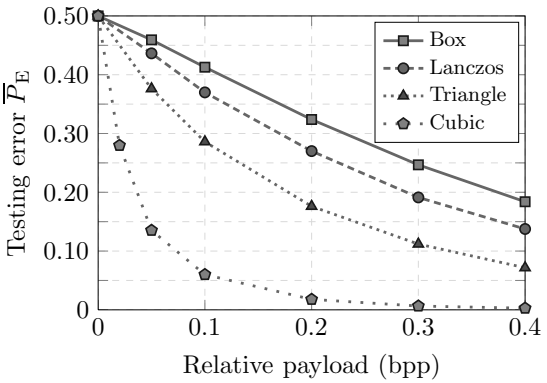

Fig. 1. Median testing error $\overline{P}_{\rm E}$ for HUGO on BOSSbase images created with four different interpolation kernels.

Statistical detectability will be measured empirically using the minimal total detection error under equal priors, $P_{\rm E} = \min_{P_{\rm FA}}(P_{\rm FA} + P_{\rm MD})/2$, where $P_{\rm FA}$ and $P_{\rm MD}$ are the false-alarm and missed-detection rates, computed from the testing set when implementing the classifier using the ensemble [4]. We denote by $\overline{P}_{\rm E}$ the median testing error over ten random splits of the database into two halves.

In this section, we represent images using the 12,753-dimensional spatial rich model SRMQ1 [5]. Figure 1 shows $\overline{P}_{\rm E}$ for HUGO (implemented with $\sigma = \gamma = 1$ and the threshold $T = 255$) for four different versions of the BOSSbase. Everything else being equal, the individual scenarios differ only in a single parameter of the `convert`'s image resizing algorithm – the interpolation kernel. The four choices were: box, Lanczos [6] (default), triangle, and cubic. The differences are rather striking. For example, at the relative payload 0.2 bpp (bits per pixel), the detection error dropped from 0.27 with the default Lanczos kernel to an almost perfect detectability (error 0.02) with bicubic interpolation. Apparently, the BOSS competition outcome (and HUGO's security) would be viewed in a completely different light depending on the BOSS organizers' choice of the interpolation algorithm.

3. FURTHER INVESTIGATION

3.1. Notation and preliminaries

Representing a two-dimensional scene (the "reality") by a function $f : \mathbb{R}^2 \to \mathbb{R}$, a digital image $\mathbf{X}$ registered by the imaging sensor can be defined as a (quantized) sampled portion of f:

$$\mathbf{X}(x,y) = Q\left(C_{\Delta,\Theta}(x,y)\cdot f(x,y)\right), \quad (1)$$

where Q is a scalar quantizer and $C_{\Delta,\Theta}(x,y)$ denotes a discrete sampling function,

$$C_{\Delta,\Theta}(x,y) = \sum_{k=0}^{M-1}\sum_{l=0}^{N-1}\delta(x-x_0-k\Delta)\delta(y-y_0-l\Delta), \quad (2)$$

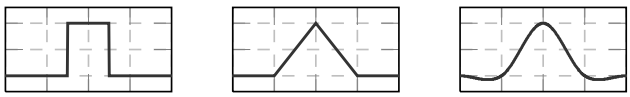

Fig. 2. The box, triangle, and cubic interpolation kernels.

with the parameter vector $\Theta = (x_0, y_0, M, N)$ and $\delta(x)$ defined as $\delta(0) = 1$ and $\delta(x) = 0\ \forall x \neq 0$. Note that $\mathbf{X}$ is non-zero only at $M \times N$ equally-spaced locations of a rectangular portion of the real scene f uniquely defined by Θ.

The actual process of image resizing is executed by means of a linear interpolation filter with a convolution kernel $\varphi : \mathbb{R}^2 \to \mathbb{R}$ satisfying $\int \varphi(x,y)\mathrm{d}x\mathrm{d}y = 1$. Formally, resizing $\mathbf{X}$ by a factor of k, one obtains:

$$\mathbf{X}^{(k)}(x,y) = Q\left(C_{\Delta/k,\Theta^{(k)}}(x,y)\cdot(\mathbf{X}*\varphi)(x,y)\right), \quad (3)$$

where the convolution $(\mathbf{X}*\varphi)(x,y)$ serves as an approximation of the reality $f(x,y)$. We allow the new parameter vector $\Theta^{(k)} = (x_0^{(k)}, y_0^{(k)}, \lfloor M/k \rfloor, \lfloor N/k \rfloor)$ to have a different starting point of the grid $(x_0^{(k)}, y_0^{(k)})$ than the original image $\mathbf{X}$, i.e., the first pixel of the resized image does not have to coincide with the first pixel of the original image. Not only does this correspond to the specific implementations in common image-processing tools, but also, as will be shown, the position of the point $(x_0^{(k)}, y_0^{(k)})$ plays a crucial role in steganalysis.

For simplicity, in the rest of the paper we assume that $M = N$, $x_0^{(k)} = y_0^{(k)}$ for all k, and $\varphi(x,y) = \varphi(x)\varphi(y)$. The variable k will *exclusively* denote the resizing factor.

3.2. Image database and the resizing algorithm

Since the BOSSbase is a collection of images from eight different cameras, its images have been resized by different factors, which makes it unsuitable for isolating the subtle effects of interpolation. Thus, for our study, we collected $3,000$ DNG images from a single camera (Leica M9) in the native resolution of $3,472 \times 5,216$ pixels, demosaicked them using UFRaw (with the setup used during the BOSS competition), and converted to 8-bit grayscale. The resulting database of $3,000$ never compressed (and not resized) images is the mother database for all our subsequent experiments.

For image resizing, we used the Matlab function `imresize` with the nearest neighbor (box), bilinear (triangle), and bicubic (cubic [7]) interpolation kernels $\varphi_{\rm b}, \varphi_{\rm t}, \varphi_{\rm c}$, respectively (see Figure 2).

To resample an image by factor k, we apply the interpolation formula (3) and then crop the resulting image $\mathbf{X}^{(k)}$ to the central 512×512 region. The cropping was included to eliminate the effects of the SRL on images of different sizes while keeping the statistical properties of pixels. By default, the value of the point $(x_0^{(k)}, y_0^{(k)})$, i.e., the location of the first

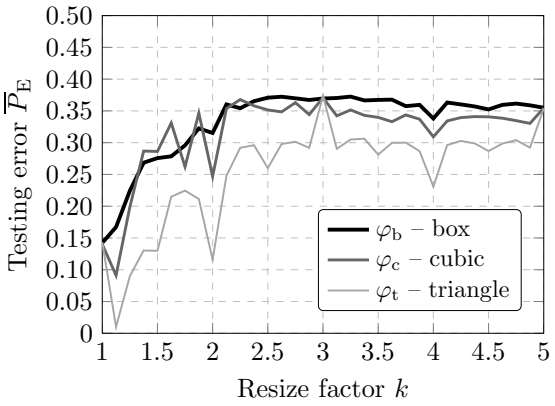

Fig. 3. Steganalysis of LSBM in images resized by different interpolation kernels.

pixel of the resized image, is calculated as

$$x_0^{(k)} = y_0^{(k)} = (k+1)/2, \tag{4}$$

which corresponds to centering the sampling points of $\mathbf{X}^{(k)}$ within the frame of the original image $\mathbf{X}$.

3.3. Steganography and steganalysis

For illustrative purposes, we attack LSB matching (LSBM). The stego images were created by changing the relative portion β of pixels (the change rate) by either increasing or decreasing their values by 1, equiprobably.

To speed up the experiments, we did not use a rich model to represent the images and, instead, opted for a more compact feature space that is adequate for the purposes of revealing the effects of resampling on security. The features are constructed from a four-dimensional co-occurrence matrix formed by four horizontally and vertically neighboring rounded noise residuals obtained by convolving the image with the kernel

$$\mathbf{K} = \begin{pmatrix} -0.25 & +0.5 & -0.25 \\ +0.5 & 0 & +0.5 \\ -0.25 & +0.5 & -0.25 \end{pmatrix}, \tag{5}$$

originally introduced in [8]. Utilizing the sign and directional symmetries of natural images [9], the co-occurrence dimensionality is reduced to 169.

3.4. Results and interpretation

Figure 3 shows the steganalysis results of LSBM at a fixed change rate achieved for a range of downsampling factors $k \in [1, 5]$ and the interpolation kernels shown in Figure 2. Note that this could be achieved by calling the Matlab's function `imresize` with kernels 'box', 'triangle', and 'cubic', and turning the antialiasing off. We remind that after resizing, all images were always cropped to the central 512×512 region to eliminate the effects of the SRL.

The testing error generally grows with increasing k, which is to be expected since downsampling decreases the dependencies among pixels. The differences between kernels are solely due to different (linear) pixel combinations created during the process of interpolation, which results in qualitatively different dependencies among pixels. These are eventually disturbed by embedding and detected by steganalysis. One of the main messages of this article is that even a small change in the interpolation kernel can significantly change the outcome of steganalysis.

The progress of the errors in Figure 3 is far from monotonous in k, which is especially apparent for the triangle kernel φ_{t}, where the sudden drops and increases sometimes exceed 5% of the error value. There are two reasons for such rapid changes – the distance between two pixels at the resolution k and the position of the first pixel in the resized image, $(x_0^{(k)}, y_0^{(k)})$. For $k = 2$, for example, equation (4) yields $x_0^{(k)} = y_0^{(k)} = 1.5$, and the pixel difference is 2. Thus, every pixel of the resized image is *exactly* in the middle between two pixels of the original image; the triangle kernel *averages* both neighboring pixels. This can be seen as a simple denoising operator that increases local correlations among pixels and thus makes steganalysis easier. Similar arguments could be applied for the case $k = 4$. The initial drop at $k \approx 1.1$ (for φ_{c} and φ_{t}) is again caused by the increased strength of dependencies among neighboring pixels – this time due to the fact that the pixel grids at both scales are misaligned and most pixels from the original resolution contribute to *two* neighboring pixels of the resized image. Note that this is not the case for the box kernel, a simple nearest neighbor interpolator, hence the missing initial drop.

Lastly, note the identical performance of all three kernels at odd factors $k = 3, 5$ (and trivially also at $k = 1$). In these cases, the pixel locations of the resized image always coincide with certain pixels from the original grid, and since all three kernels vanish at integer values, they are essentially identical. Furthermore, for the triangle kernel, this means a sudden decrease in neighboring pixel correlations (compared to the close, non-integer values of k) and thus a sudden drop in steganalysis performance (increased error).

The presented qualitative interpretation of Figure 3 provides insight into the inner workings of the image interpolation and the importance of its individual components for steganalysis (e.g., the alignment of the resized grid). In the next section, we approach the problem from a theoretical perspective that points to an intriguing new scaling law.

4. SCALING W.R.T. IMAGE RESOLUTION

4.1. Image model

In this section we study the effects of image resizing on steganalysis analytically. To this end, we restrict ourselves to the simplest box kernel and assume that pixels in rows/columns

of images are first-order Markov chains (MCs) over the space of grayscales and are thus fully characterized by a transition probability matrix (TPM) $\mathbf{A} = (a_{ij})$. Following [10], we adopt the exponential cover model,

$$a_{ij} = 1/Z_i \exp(-(|i-j|/\tau))^\gamma, \tag{6}$$

whose parameters can be estimated from a large number of images. The parameters τ and γ were estimated using the method of moments [11] from the Leica M9 database.

Resizing images by a factor $k \geq 1$, $k \in \mathbb{R}$, changes the TPM $\mathbf{A} \to \mathbf{A}^k$, where the matrix power is defined in a general sense and can be evaluated via eigenvalue decomposition for non-integer k. This allows us to study the scaling effects of image resizing on steganographic security solely based on the statistical properties of the original cover source (*non-resized* images).

4.2. Scaling factor $\alpha(k)$

It is well-known that the leading term of the KL divergence between cover and stego objects is quadratic in the change rate β:

$$D(k;\beta) = \frac{1}{2} n \beta^2 I(k), \tag{7}$$

where n is the cover size and $I(k)$ is the steganographic Fisher information (FI) rate for the resize factor k. For a fixed cover source described by its TPM and a fixed steganographic algorithm (LSBM in our case), the authors of [10] derived a closed-form expression for $I(k)$ (see Thm. 2 in [10]), from which one can obtain $D(k;\beta)$ at different resolutions (as a function of $\mathbf{A}^k$).

To obtain a constant level of statistical detectability (KL divergence D) after resizing by factor k, the change rate β needs to be scaled by $\alpha(k)$: $D(1;\beta) = D(k;\alpha(k)\beta)$. Since we always central-crop the image after resizing to the same size n (and thus eliminate the effect of the SRL) it is easy to see that $\alpha(k) = \sqrt{I(1)/I(k)}$. In Figure 4 (left), we show the computed values of the parameter $\alpha(k)$ for a range of scaling factors $k \in [1,3]$. These theoretically obtained results were verified in the following manner. For a fixed change rate β, we first steganalyzed LSBM using the feature vector described in Sec. 3.3, obtaining thus the median testing error $\overline{P}_{\mathrm{E}}(1,\beta)$. Next, according to the theory, the same error rate, $\overline{P}_{\mathrm{E}}(k,\alpha(k)\beta)$, should be obtained after resizing the image by the factor k and modifying the change rate β to $\alpha(k)\beta$, provided the images at both resolutions are cropped to the same dimensions (otherwise, another change rate adjustment due to the SRL would be needed). In Figure 4 (right), we show the results for several values of β for $k = 2$. Despite the simplicity of the cover model (6) and the fact that we use a machine-learning based detector, an excellent match between theory and practice is observed. We plan to include more experimental results in a journal version of this paper.

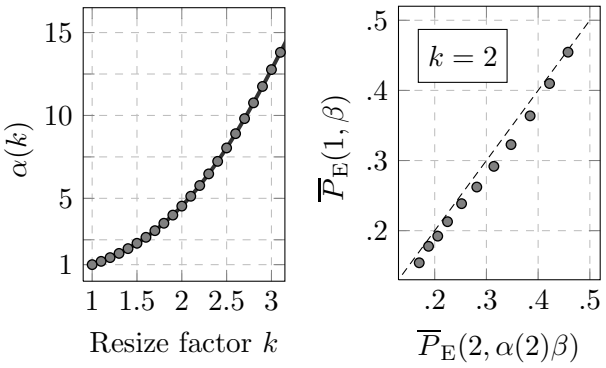

Fig. 4. Left: Derived scaling factor $\alpha(k) = \sqrt{I(1)/I(k)}$; Right: testing error $\overline{P}_{\mathrm{E}}(2,\alpha(2)\beta)$ on images resized by factor $k = 2$ embedded with change rate $\alpha(2)\beta$ vs. $\overline{P}_{\mathrm{E}}(1,\beta)$ for non-resized images ($k = 1$) embedded with change rate β for the following range of values of $\beta = \{0.01, 0.02, \ldots, 0.11\}$ ($\alpha(2) \approx 4.53$). See the text for more details.

5. CONCLUSIONS AND RELATION TO PRIOR ART

To the best knowledge of the authors the role of the image resizing algorithm and its influence on steganography and steganalysis has not been studied so far. This paper is the first step in this direction, and it reveals a surprising sensitivity of steganalysis to the choice of the interpolation kernel, as well as the exact position of the first pixel of the resized image, which affects the alignment of both pixel grids and consequently the strength of the correlations among individual pixels.

Resizing an image is a very different operation from cropping or concatenating images because the statistical properties of pixels change. Thus, the standard scaling of secure payload as stated by the square root law is no longer valid. To obtain a better understanding of this phenomenon, we theoretically analyzed the simplest resizing algorithm with the box kernel (the nearest neighbor interpolation). Adopting a Markov chain model for image pixels, we were able to compute the steganographic Fisher information rate, which allowed us to derive the scaling of the secure payload that should lead to constant statistical detectability under resizing.

Studying the impacts of image resizing on statistical properties of pixels is significant for practitioners because steganography and steganalysis are nowadays commonly benchmarked on image sets obtained by resizing full-resolution images (e.g., the BOSSbase).

The authors are preparing an expanded journal version of this article that will include the study of the effects of resizing with antialiasing as well as a more detailed study of the scaling law.

6. REFERENCES

[1] T. Filler, A. D. Ker, and J. Fridrich, "The Square Root Law of steganographic capacity for Markov covers," in *Proceedings SPIE, Electronic Imaging, Security and Forensics of Multimedia XI*, N. D. Memon, E. J. Delp, P. W. Wong, and J. Dittmann, Eds., San Jose, CA, January 18–21, 2009, vol. 7254, pp. 08 1–08 11.

[2] P. Bas, T. Filler, and T. Pevný, "Break our steganographic system – the ins and outs of organizing BOSS," in *Information Hiding, 13th International Workshop*, T. Filler, T. Pevný, A. Ker, and S. Craver, Eds., Prague, Czech Republic, May 18–20, 2011, vol. 6958 of Lecture Notes in Computer Science, pp. 59–70.

[3] T. Pevný, T. Filler, and P. Bas, "Using high-dimensional image models to perform highly undetectable steganography," in *Information Hiding, 12th International Workshop*, R. Böhme and R. Safavi-Naini, Eds., Calgary, Canada, June 28–30, 2010, vol. 6387 of Lecture Notes in Computer Science, pp. 161–177, Springer-Verlag, New York.

[4] J. Kodovský, J. Fridrich, and V. Holub, "Ensemble classifiers for steganalysis of digital media," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 2, pp. 432–444, April 2012.

[5] J. Fridrich and J. Kodovský, "Rich models for steganalysis of digital images," *IEEE Transactions on Information Forensics and Security*, vol. 7, no. 3, pp. 868–882, June 2012.

[6] A. S. Glasser, *Graphics Gems*, Morgan Kaufman, 1990.

[7] R. G. Keys, "Cubic convolution interpolation for digital image processing," *IEEE Transactions on Acoustics, Speech, and Signal Processing*, vol. ASSP-29, no. 6, pp. 1153–1160, December 1981.

[8] A. D. Ker and R. Böhme, "Revisiting weighted stego-image steganalysis," in *Proceedings SPIE, Electronic Imaging, Security, Forensics, Steganography, and Watermarking of Multimedia Contents X*, E. J. Delp, P. W. Wong, J. Dittmann, and N. D. Memon, Eds., San Jose, CA, January 27–31, 2008, vol. 6819, pp. 5 1–5 17.

[9] J. Fridrich and J. Kodovský, "Steganalysis of LSB replacement using parity-aware features," in *Information Hiding, 14th International Workshop*, Berkeley, CA, May 15–18, 2012, Lecture Notes in Computer Science, To appear.

[10] T. Filler and J. Fridrich, "Fisher information determines capacity of ϵ-secure steganography," in *Information Hiding, 11th International Workshop*, S. Katzenbeisser and A.-R. Sadeghi, Eds., Darmstadt, Germany, June 7–10, 2009, vol. 5806 of Lecture Notes in Computer Science, pp. 31–47, Springer-Verlag, New York.

[11] S. Meignen and H. Meignen, "On the modeling of DCT and subband image data for compression," *IEEE Transactions on Image Processing*, vol. 4, no. 2, pp. 186–193, February 1995.